

WinFS Research Talk

Quentin Clark

November 2003



WinFS

Agenda

- Why Longhorn, and why a storage effort
- What WinFS is
- What will make this attempt successful



Why We All Do Not Work in Cupertino



- Closed system
- No other hardware vendors
- Difficult developer platform
- Applications bundled (MacWrite, MacPaint)



Why We All Do Not Work in Cupertino



- Closed system
- No other hardware vendors
- Difficult developer platform
- Applications bundled (MacWrite, MacPaint)



IBM-Compatible PCs




- IBM, Compaq, Tandy, Dell (1984), Gateway (1985), etc.
- Healthy ecosystem, many hardware and software vendors
- Microsoft is all about developers



It's About the Platform

Each Platform Wave has produced enormous results



The Platform is under Attack

EIGHT COMPANIES FORM LINUX ALLIANCE

LINUX IN YOUR LIVING ROOM

Linux emerges as cell phone dark horse

IBM touts Linux for China

Linux nears Windows XP usability

LINUX USERS HAVE IBM ON THEIR SIDE

Massachusetts Officially Adopts Open-Source Policy

HP gives Linux a boost

HP, Red Flag partner on corporate Linux

Motorola unveils first Linux smart phone

Walmart.com to sell more Linux software

Dell tightens Linux ties

Korea launches a switch to open source

China blocks foreign software



SJMN, InfoWorld, CNET past few months headlines



The Platform is under Attack

EIGHT COMPANIES FORM LINUX ALLIANCE

LINUX IN YOUR LIVING ROOM

Linux emerges as cell phone dark horse

IBM touts Linux for China

Linux nears Windows XP usability

LINUX USERS HAVE IBM ON THEIR SIDE

Massachusetts Officially Adopts Open-Source Policy

HP gives Linux a boost

HP, Red Flag partner on corporate Linux

Motorola unveils first Linux smart phone

Walmart.com to sell more Linux software

Dell tightens Linux ties

Korea launches a switch to open source

China blocks foreign software



SJMN, InfoWorld, CNET past few months headlines



The Platform is under Attack

EIGHT COMPANIES FORM LINUX ALLIANCE

LINUX IN YOUR LIVING ROOM

Linux emerges as cell phone dark horse

IBM touts Linux for China

Linux nears Windows XP usability

LINUX USERS HAVE IBM ON THEIR SIDE

Massachusetts Officially Adopts Open-Source Policy

HP gives Linux a boost

HP, Red Flag partner on corporate Linux

Motorola unveils first Linux smart phone

Walmart.com to sell more Linux software

Dell tightens Linux ties

Korea launches a switch to open source

China blocks foreign software



SJMN, InfoWorld, CNET past few months headlines



The Platform is under Attack

EIGHT COMPANIES FORM LINUX ALLIANCE

LINUX IN YOUR LIVING ROOM

Linux emerges as cell phone dark horse

IBM touts Linux for China

Linux nears Windows XP usability

LINUX USERS HAVE IBM ON THEIR SIDE

Massachusetts Officially Adopts Open-Source Policy

HP gives Linux a boost

HP, Red Flag partner on corporate Linux

Motorola unveils first Linux smart phone

Walmart.com to sell more Linux software

Dell tightens Linux ties

Korea launches a switch to open source

China blocks foreign software



SJMN, InfoWorld, CNET past few months headlines



The Platform is under Attack

EIGHT
LINU
Linux
IBM t
Linux
LINU
Massa
HP gives Linux a boost



UNIX ALL
DM
ark horse

ity
THEIR S
Open-So



on corporate Linux
Linux smart phone
more Linux software




SJMN, InfoWorld, CNET past few months headlines




It's About the Platform

Each Platform Wave has produced enormous results



It's About the Platform

Each Platform Wave has produced enormous results



The Next Platform – Longhorn

- **Industry platform for the next wave**
 - A next-gen desktop OS platform, new applications and services
 - DOS → Win16 → Win32 → LAPI
 - Built in next-gen framework (managed code)
- **Provide new innovation opportunities to ISVs**
 - Presentation (Avalon, Aero)
 - Communication (Indigo)
 - Storage (WinFS, Aero)
- **It's about the industry**
 - Innovation
 - Productivity



Why Focus on Integrated Storage?

- Digital-born data
 - Business processes generating more digital data
 - Digital photography and video
- Search
 - Knowledge workers spend 15-25% of time searching (AIIM)
 - KW's spend 2.5 hours per day searching for information (IDC)
- Fortunately, technology supports these trends
 - 70%/year individual disk capacity growth over last decade



Customer Benefits

- Find
 - Information is organized in the way I think about it
 - I can find my data
 - I have one view of my information, no matter what app I used to create it
 - I can easily manipulate my data from any application
- Relate
 - I think about information not files
 - My data is smarter and related together
- Act
 - I'm not overloaded with irrelevant information
 - My information, my way
 - I can work anywhere, anytime, anyplace (including offline)



WinFS Developer Benefits

- New Filesystem Capabilities
 - Developer benefits even without changing file formats by leveraging data model, relationships, promotion
 - My current applications just continue to work on Longhorn, and new capabilities
- New Windows primitives
 - New things in the os to integrate with (contacts, calendar, media, locations)
 - I can take advantage of a simple pre-built structures that I can extend
- New data subsystem capabilities
 - I can easily build a smart connected application
 - I can easily and securely get to data shared by other apps
 - I can easily expose my branded data to other applications and the shell



A Few Scenarios

- Contacts
 - Applications can use the Contact APIs, and not have to understand the myriad vertical PIM solutions to provide contact integration
- Documents
 - “The truth is out there”
 - Document solutions built on WinFS can leverage the richness of the local store and peer group synchronization
 - Annotations on documents and other items
 - Finding all types items that relate to specific work threads
 - Following relationships
 - Accessing the file with existing Windows applications
- “Personal Memories”
 - Relating people, photos, video, calendar events, and places opens a vast set of capabilities
 - User-driven categorization of their data by navigation and relationships
 - Applications can build new types and relationships to enrich the end-user experience



So, what is WinFS?

WinFS is the active storage platform for organizing, searching and sharing data

- Where all end-user data lives in Longhorn
- A secure place to store data
- A data model built on database technology and the filesystem
- A set of services that make your data active
- A set of domain-specific schemas
 - Contacts, Documents, Pictures, Tasks, Audio, Events, etc.
- Part of the Longhorn developer innovation push
- Married to a new Shell UX
- Really Hard



WinFS Is *Not*

- Abandoning legacy applications
- Forcing the world to re-write everything
- Done
 - Design complete by end of year ("complete the innovation")
 - Fundamental model complete end of year
 - Lots of work to do next year even (perf, scale, basics, server plan, self-management, etc.)
- Ignorant of the mistakes of projects past
 - Not building from scratch
 - Key technical decision (e.g. let NTFS deal with file handles)
- Going to be successful without our partners (Shell, Outlook12, STS, tools, etc.)
- The thing that will make Longhorn slip into the next decade ☺



WinFS Item Store

Complete data APIs

- Object, Relational, and XML

Schema framework

- Extensible "everyday info" schemas
- ISV-build schema

Active data

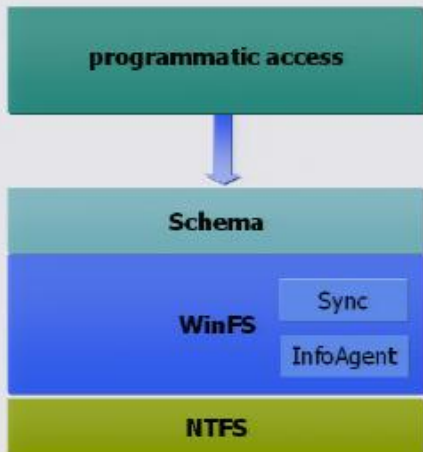
- Synchronization
InfoAgent

Structured "Item" store

- Data model
- Files/Folders → Items/Relationships

File System

- NTFS Semantics for Filestreams
- Backwards Compatible with Win32



WinFS Data Model

Items

- WinFS Items are the primary things that applications work on
- Can be related to other items
- Move, copy, delete, serialize/de-serialize

Relationships

- Provide annotation for relating items together
- Schema can model complex items
- Containment, reference, embedding, categories, etc

Searching / Organization

- Explicit folders, virtual folders, implicit through queries, explicit connections between items, etc


pictures



documents



people



music



annotations



events



e-mail



Synchronization Overview

Scenarios

- Synchronize one WinFS with another
 - Keep My Contacts and My Files in sync across my home machines
 - Peer to Peer sharing
- Synchronize WinFS with other data sources
 - Keep My Contacts in sync with online email contacts, enterprise CRM, etc.

Approach

- Multi-master replication (peer-to-peer)
 - Replicas make changes independently
- Folders are unit of scope, items are unit of consistency
- Net-change synchronization (as opposed to action-)
 - Looking at cumulative changes, not logs
- A set of common services for all data sources and all schemas
 - Change tracking, change enumeration, conflict handling, etc.



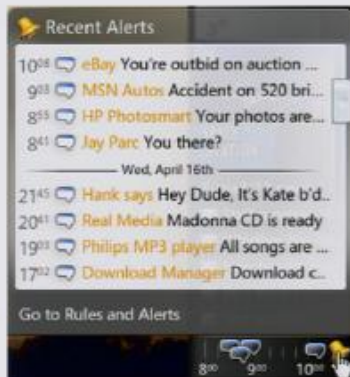
WinFS Notifications

'Active Data' – Subscribe to WinFS changes

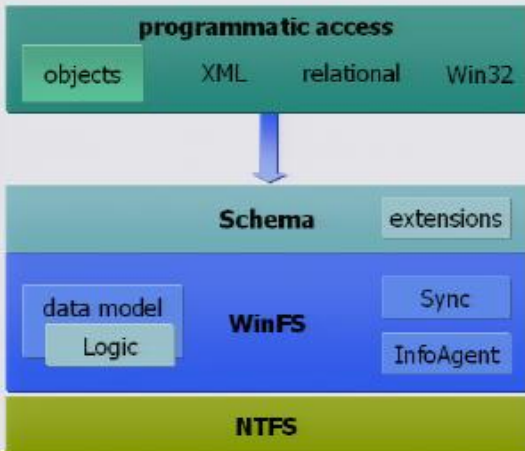
- Item change subscriptions
- Item Domain containment / query subscriptions

InfoAgent Integration

- Inclusive set of events, contexts, and actions
- Preferences stored as WinFS items
- Unified management of notification rules



WinFS for ISVs



1 Use the WinFS API

2 Extend Windows types

3 Create new types

4 Leverage InfoAgent

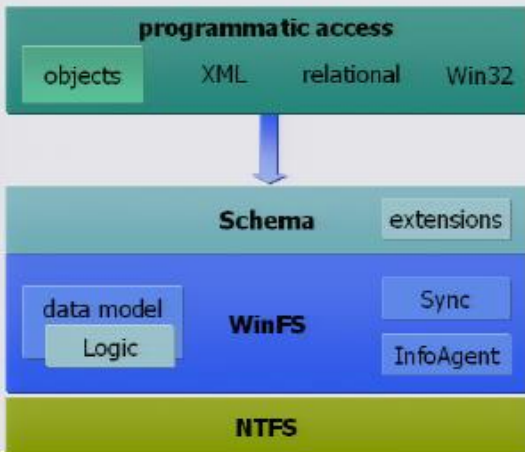
4 Use Sync, build Adaptors

5 Add in-store Logic

5 Write file promoters



WinFS for ISVs



1 Use the WinFS API

2 Extend Windows types

3 Create new types

4 Leverage InfoAgent

4 Use Sync, build Adaptors

5 Add in-store Logic

5 Write file promoters



This Time it's Different (no, really)

- We are looking at the whole system
 - Shell, Outlook, STS, Dev tools
 - Double-edged sword
- We don't have to invent the whole world
 - More supporting technology
 - CLR, NUI, IS system, FullText, NTFS
- Higher level of commitment
 - Not without it's moments ;), but the execs are bought-in
- Key technical decisions
 - Don't try to just re-implement an old system (MAPI, file handles)
- We have relational nerds
 - The heart of solving Integrated Storage is a relational model



This Time it's Different

- We are looking at the whole system
 - Shell, Outlook, STS, Dev tools
 - Double-edged sword
- We don't have to invent the whole world
 - More supporting technology
 - CLR, NUI, IS system, FullText, NTFS
- Higher level of commitment
 - Not without it's moments ;), but the execs are bought-in
- Key technical decisions
 - Don't try to just re-implement an old system (MAPI, file handles)
- We have relational nerds
 - The heart of solving Integrated Storage is a relational model



This Time it's Different

- We are looking at the whole system
 - Shell, Outlook, STS, Dev tools
 - Double-edged sword
- We don't have to invent the whole world
 - More supporting technology
 - CLR, NUI, IS system, FullText, NTFS
- Higher level of commitment
 - Not without it's moments ;), but the execs are bought-in
- Key technical decisions
 - Don't try to just re-implement an old system (MAPI, file I/O)
- We have relational nerds
 - The heart of solving Integrated Storage is a relational model



What's so Cool?

- Revolutionize everyday information interactions
 - Completely new user paradigm, proving much more intuitive
 - Consumers use the PC differently than ever before
 - Don't you want something better than FindFirst FindNext?
- Drive innovation
 - What new applications that were not possible before will be built?
 - PCs come back into the limelight - "Personal Memories" vision
- Drive the stock price (up ;))
 - Knowledge Workers spent 2.5 hours/day looking for information*
 - Productivity gains drive economic and company growth
 - ISV virtuous cycle

* IDC

